PATENT
Docket No. **0023-0118**

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

In re Patent Application of	)	Mail Stop: **Appeal Brief - Patents**
Der-Hwa GAN et al.	)	Group Art Unit: 2664
Application No.: 09/354,640	)	Examiner: C. Ho
Filed: July 15, 1999	)	
For: METHOD AND APPARATUS FOR FAST REROUTE IN A CONNECTION-ORIENTED NETWORK	)	

U.S. Patent and Trademark Office
Customer Window, Mail Stop **Appeal Brief - Patents**
Randolph Building
401 Dulany Street
Alexandria, Virginia 22314

APPEAL BRIEF

This Appeal Brief is submitted in response to the final Office Action, dated January 9, 2006, and in support of the Notice of Appeal, filed March 8, 2006.

I. REAL PARTY IN INTEREST

The real party in interest in this appeal is Juniper Networks, Inc.

II. RELATED APPEALS, INTERFERENCES, AND JUDICIAL PROCEEDINGS

Appellants are unaware of any related appeals, interferences or judicial proceedings.

III. STATUS OF CLAIMS

Claims 1-4, 6, 8-21, and 24 are pending in this application.

Claims 1-4, 6, 8-21, and 24 were finally rejected in the Office Action, dated January 9, 2006, and are the subject of the present appeal. These claims are reproduced in the Claim Appendix of this Appeal Brief.

IV. STATUS OF AMENDMENTS

No amendment was filed subsequent to the final Office Action, dated January 9, 2006.

V. SUMMARY OF CLAIMED SUBJECT MATTER

In the paragraphs that follow, each of the independent claims that is involved in this appeal and each dependent claim that is argued separately will be recited followed in parenthesis by examples of where support can be found in the specification and drawings.

Claim 1 recites a network (140, Fig. 1a) for forwarding packets from a source device (100, Fig. 1a) to a destination device (110, Fig. 1a), said network including a plurality of network elements including a plurality of nodes (130, Fig. 1a) and connecting links (150, Fig. 1a), the plurality of nodes including at least one alternative-route-enabled node (201, Fig. 2a) and at least one non-alternative-route-enabled node (203, Fig. 2a), wherein the at least one non-alternative-route-enabled node comprises a storage space to store an initial route from the source device to

the destination device (Fig. 6b; pg. 11, lines 16-20; pg. 8, lines 11-12); a mechanism to detect failure in a downstream network element in the initial route (315, Fig. 3a; pg. 8, line 27, to pg. 9, line 2); and a forwarder to automatically forward a failure message upstream along the initial route to an alternative-route-enabled node, the failure message causing the alternative-route-enabled node to begin forwarding packets on an alternative route (322, 324, Fig. 3a; pg. 9, lines 2-9).

Claim 8 recites a method for forwarding packets from a source device (100, Fig. 1a) to a destination device (110, Fig. 1a) in a network of interconnected elements including nodes (130, Fig. 1a) and links (150, Fig. 1a), comprising determining an initial route, the initial route including at least one alternative-route enabled node and at least one non-alternative-route-enabled node, the at least one alternative-route-enabled node and the at least one non-alternative-route-enabled node storing an initial route from the source device to the destination device (330, Fig. 3b; pg. 8, lines 4-12); determining an alternative route by identifying the at least one alternative-route-enabled node in the initial route, identifying downstream interconnected elements, and generating the alternative route based on the identified at least one alternative-route-enabled node and the identified downstream interconnected elements (308, Fig. 3a; 505-525, Fig. 5a; pg. 8, lines 19-24, pg. 10, lines 7-15); forwarding packets on the initial route (310, Fig. 3a; pg. 8, lines 24-25); detecting a failed element (315, Fig. 3a; pg. 8, lines 27-30); and automatically forwarding packets on the alternative route without communicating with either the source or the destination (322, 320, 324, Fig. 3a; pg. 8, line 30, to pg. 9, line 9).

Claim 9 recites that the determining the initial route further comprises determining a short path from the destination device to the source device within the network (410, Fig. 4; pg. 9, lines

14-19); refining the path according to administrative constraints (415, 420, Fig. 4; pg. 9, lines 19-29); and establishing the path as the initial route (425, Fig. 4; pg. 9, lines 29-30).

Claim 10 recites that the refining the path comprises rejecting the path exceeding bandwidth allocation and hop limit (pg. 9, lines 22-30).

Claim 11 recites that the determining the alternative route further comprises determining a shortest route from a node preceding the failed element to the destination device within the network (pg. 8, lines 13-14; pg. 9, lines 29-30); refining the route to exclude the failed element on the initial route (520, Fig. 5a; pg. 10, lines 12-14); and establishing the alternative route for forwarding packets (522, Fig. 5a; pg. 10, lines 14-15).

Claim 13 recites that the determining the alternative route comprises reserving bandwidth available on the initial route (pg. 4, lines 13-21; pg. 10, lines 7-30); generating the alternative route by invoking a routing protocol (pg. 4, lines 13-21; pg. 10, lines 14-15); refining the alternative route by excluding the failed element (pg. 4, lines 13-21; pg. 10, lines 11-14); and establishing the alternative route (pg. 4, lines 13-21; pg. 10, lines 14-15).

Claim 14 recites a method for forwarding packets from a source device (100, Fig. 1a) to a destination device (110, Fig. 1a) in a network of interconnected elements including nodes (130, Fig. 1a) and links (150, Fig. 1a), comprising determining an initial route by determining a short path from the destination device to the source device within the network (305, Fig. 3b; 410, Fig. 4; pg. 9, lines 15-17), refining the path according to administrative constraints (415, 420, Fig. 4; pg. 9, lines 19-29), and establishing the path as the initial route (425, Fig. 4; pg. 9, lines 29-30), the initial route being prioritized to establish a hierarchy for preemption in routing network traffic (pg. 10, lines 1-6); determining an alternative route (308, Fig. 3a; pg. 8, lines 19-21);

forwarding packets on the initial route (310, Fig. 3a; pg. 8, lines 24-25); detecting a failed element (315, Fig. 3a; pg. 8, lines 27-29); and automatically forwarding packets on the alternative route without communicating with either the source or the destination (322, 320, 324, Fig. 3a; pg. 8, line 30, to pg, 9, line 9).

Claim 15 recites that the determining the alternative route comprises checking bandwidth allocation (540, 545, Fig. 5; pg. 4, lines 13-21; pg. 8, lines 7-18; pg. 10, line 22, to pg. 11, line 6).

Claim 16 recites that the checking bandwidth allocation comprises dynamically balancing capacity of nodes and links (540, 545, Fig. 5; pg. 4, lines 13-21; pg. 8, lines 7-18; pg. 10, line 22, to pg. 11, line 6).

Claim 17 recites that the determining the alternative route comprises reserving bandwidth available on the initial route (pg. 4, lines 13-21; pg. 10, lines 7-30); identifying a plurality of nodes associated with the failed element according to network configuration information (pg. 4, lines 13-21; pg. 10, line 11); generating the alternative route excluding the failed element and the plurality of nodes (pg. 4, lines 13-21; pg. 10, lines 11-14); and establishing the alternative route (pg. 4, lines 13-21; pg. 10, lines 14-15).

Claim 18 recites a method for locally rerouting packets traveling on an established route when a node in a network of interconnected nodes fails, the method comprising computing, at select intermediary nodes along the established route, an alternative route leading from the select intermediary node to the destination device of the established route (308, Fig. 3a; pg. 8, lines 19-21); storing, at each of the select intermediary nodes, the alternative route (pg. 8, lines 22-24); determining locally that the established route has failed (315, Fig. 3a; pg. 8, lines 27-29); and automatically forwarding packets on the alternative route (322, 320, Fig. 3a; pg. 8, line 30, to pg.

9, line 9).

Claim 19 recites that computing the alternative route comprises reserving bandwidth available on the established route (pg. 4, lines 13-21; pg. 10, lines 7-30); identifying a plurality of nodes associated with the failed node according to network configuration information (pg. 4, lines 13-21; pg. 10, line 11); generating the alternative route excluding the failed node and the plurality of nodes (pg. 4, lines 13-21; pg. 10, lines 11-14); and establishing the alternative route (pg. 4, lines 13-21; pg. 10, lines 14-15).

Claim 20 recites that computing the alternative route further comprises locating a set of established routes with a same destination device and same administrative constraints as the established route (540, 545, Fig. 5b; pg. 11, lines 1-4); finding a common node, downstream from the failed node, after which the set of established routes and the established route utilize the same network elements (550, Fig. 5b; pg. 11, lines 4-5); establishing a new route from the common node to the destination device (555, Fig. 5; pg. 11, lines 5-6); and incorporating the new route into the alternative route (pg. 11, lines 5-6).

Claim 24 recites a network for forwarding packets from a source device (100, Fig. 1a) to a destination device (110, Fig. 1a) and including a plurality of intermediate network nodes (130, Fig. 1a), the plurality of intermediate network nodes comprising at least one first node (201, Fig. 2a) configured to store an initial route from the source device to the destination device and at least one alternative route from the source device to the destination device (pg. 8, lines 11-12 and 19-24), detect a failure in a downstream network node in the initial route (315, Fig. 3a; pg. 8, lines 27-29), and automatically forward a packet to a node on one of the at least one alternative route in response to detecting the failure (322, 320, Fig. 3a; pg. 8, line 30, to pg. 9, line 9); and at

least one second node (203, Fig. 2a) configured to store the initial route (pg. 8, lines 11-12), detect a failure in a downstream network node in the initial route (315, Fig. 3a; pg. 8, lines 27-29, and forward a failure message to an upstream first node in response to detecting the failure, the failure message causing the upstream first node to automatically forward a packet to a node on one of the at least one alternative route (322, 324, Fig. 3a; pg. 8, line 30, to pg. 9, line 9).

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

A. Claims 1-4, 6, 8, 9, 11, 12, 14, 18, 21, and 24 stand rejected under 35 U.S.C. § 103(a) as unpatentable over Haskin et al. (U.S. Patent No. 6,813,242) in view of McAllister et al. (U.S. Patent No. 6,697,329).

B. Claims 10, 13, 15-17, and 19-21 stand rejected under 35 U.S.C. § 103(a) as unpatentable over Haskin et al. (U.S. Patent No. 6,813,242) in view of McAllister et al. (U.S. Patent No. 6,697,329), and further in view of Rexford (U.S. Patent No. 6,633,544).

VII. ARGUMENTS

A. The rejection under 35 U.S.C. § 103(a) based on Haskin et al. (U.S. Patent No. 6,813,242) and McAllister et al. (U.S. Patent No. 6,697,329) should be reversed.

The initial burden of establishing a *prima facie* basis to deny patentability to a claimed invention always rests upon the Examiner. In re Oetiker, 977 F.2d 1443, 24 USPQ2d 1443 (Fed. Cir. 1992). In rejecting a claim under 35 U.S.C. § 103, the Examiner must provide a factual basis to support the conclusion of obviousness. In re Warner, 379 F.2d 1011, 154 USPQ 173 (CCPA 1967). Based upon the objective evidence of record, the Examiner is required to make

the factual inquiries mandated by Graham v. John Deere Co., 86 S.Ct. 684, 383 U.S. 1, 148 USPQ 459 (1966). The Examiner is also required to explain how and why one having ordinary skill in the art would have been realistically motivated to modify an applied reference and/or combine applied references to arrive at the claimed invention. Uniroyal, Inc. v. Rudkin-Wiley Corp., 837 F.2d 1044, 5 USPQ2d 1434 (Fed. Cir. 1988).

In establishing the requisite motivation, it has been consistently held that the requisite motivation to support the conclusion of obviousness is not an abstract concept, but must stem from the prior art as a whole to impel one having ordinary skill in the art to modify a reference or to combine references with a reasonable expectation of successfully achieving some particular realistic objective. See, for example, Interconnect Planning Corp. v. Feil, 227 USPQ 543 (Fed. Cir. 1985). Consistent legal precedent admonishes against the indiscriminate combination of prior art references. Carella v. Starlight Archery, 804 F.2d 135, 231 USPQ 644 (Fed. Cir. 1986); Ashland Oil, Inc. v. Delta Resins & Refractories, Inc., 776 F.2d 281, 227 USPQ 657 (Fed. Cir. 1985).

1. Claims 1-4 and 6.

Claim 1 is directed to a network for forwarding packets from a source device to a destination device, where the network includes a plurality of network elements including a plurality of nodes and connecting links. The plurality of nodes includes at least one alternative-route-enabled node and at least one non-alternative-route-enabled node. The at least one non-alternative-route-enabled node comprises a storage space to store an initial route from the source device to the destination device; a mechanism to detect failure in a downstream network element in the initial route; and a forwarder to automatically forward a failure message upstream along

the initial route to an alternative-route-enabled node, where the failure message causes the alternative-route-enabled node to begin forwarding packets on an alternative route. Haskin et al. and McAllister et al., whether taken alone or in any reasonable combination, do not disclose or suggest this combination of features.

For example, Haskin et al. and McAllister et al. do not disclose or suggest a plurality of nodes including at least one alternative-route-enabled node and at least one non-alternative-route-enabled node. The Examiner relies on element 1 in Fig. 1 of Haskin et al. as allegedly corresponding to at least one alternative-route-enabled node and on element 3 or 5 in Fig. 1 of Haskin et al. as allegedly corresponding to at least one non-alternative-route-enabled node (final Office Action, pg. 4). Appellants respectfully disagree with the Examiner's interpretation of the disclosure of Haskin et al.

Element 1 in Fig. 1 of Haskin et al. corresponds to a switch. As clearly illustrated in Fig. 1, switch 1 has the ability to transfer data over a primary path 13 and an alternative path 12. Thus, switch 1 would be construed as an alternative-route-enabled node. Elements 3 and 5 in Fig. 1 of Haskin et al. also correspond to switches. As clearly illustrated in Fig. 1, switches 3 and 5 have the ability to transfer data over a primary path 35 and 57, respectively, and an alternative path 34 and 56, respectively. Thus, switches 3 and 5 would also be construed as alternative-route-enabled nodes. The Examiner does not explain why one skilled in the art could reasonably construe switches 3 and 5 to be non-alternative-route-enabled nodes when Haskin et al. specifically discloses that alternative routes are provided for switches 3 and 5 (see, for example, col. 3, lines 4-9).

While the Examiner appears to rely on Haskin et al. for allegedly disclosing a non-

alternative-route-enabled node, the Examiner also appears to admit that Haskin et al. does not disclose this feature by relying on McAllister et al. for allegedly disclosing a non-alternative-route-enabled node (see, for example, final Office Action, pp. 2 and 5). In particular, the Examiner relies on element 32A in McAllister et al. as allegedly corresponding to an alternative-route-enabled node, on element 32B in McAllister et al. as allegedly corresponding to a non-alternative-route-enabled node (final Office Action, pp. 2 and 5). Appellants respectfully disagree with the Examiner's interpretation of McAllister et al.

Element 32A in McAllister et al. corresponds to an ingress node (see, for example, col. 8, lines 26-31) and element 32B corresponds to a node, which is attached to ingress node 32A (see, for example, Figs. 2A and 2B, and col. 6, lines 10-12). McAllister et al. in no way discloses or suggests that node 32B is a non-alternative-route-enabled node. In stark contrast, McAllister et al. specifically discloses that node 32B may transmit data via two separate paths (see, for example, col. 7, lines 35-49, where McAllister et al. specifically discloses that operator directed route (ODR) soft permanent virtual circuits (SPVCs) can be established through a primary route including node 32A; node 32B, line 36BC; node 32C; and 32D, and an alternate route including node 32A; node 32B, line 36BC'; node 32C; and 32D). Thus, node 32B is capable of transmitting data via two separate routes (a primary route and an alternative route). Therefore, node 32B would be construed as an alternative-route-enabled node and not, as alleged by the Examiner, a non-alternative-route-enabled node.

Even assuming, for the sake of argument, that one skilled in the art at the time of Appellants' invention could reasonably construe McAllister et al.'s node 32B as a non-alternative-route-enabled node (a point that Appellants do not concede), Appellants submit that McAllister

et al. does not disclose or suggest that node 32B includes a storage space to store an initial route from a source device to a destination device, as required by claim 1. The Examiner relies on cols. 45-57, of McAllister et al. for allegedly disclosing a non-alternative-route-enabled node that includes a storage space to store an initial route from a source device to a destination device. At the outset, Appellants note that the McAllister et al. document does not include columns numbered 45-57. Appellants assume that this is a typographical error.

Nonetheless, McAllister et al.'s Fig. 5 depicts a relational database structure 52 that is used by a network node 32 to keep track of ODR SPVCs managed by that network node 32 (col. 13, lines 35-37). McAllister et al. appears to disclose that ingress node 32A manages the ODR SPVCs of network 30 (see, for example, col. 9, lines 6-21, where McAllister et al. specifically discloses that ingress node 32A stores the configuration data for and initiates the ODR SPVC and that the nature of the ODR SPVC is transparent to the remainder of network 30). The Examiner has not pointed to any section of McAllister et al. that discloses or suggests that node 32B includes a storage space to store an initial route from a source device to a destination device, as required by claim 1.

Moreover, McAllister et al. does not disclose or suggest that relational database structure 52 stores an initial route from a source device to a destination device, as required by claim 1. Instead, McAllister et al. specifically discloses that relational database structure 52 stores an indexed table 54 of P-NNI network node identifiers, and an indexed table 56 of compressed ODR SPVC designated transit lists (DTLs) and an ODR SPVC list 58. McAllister et al. does not disclose or suggest that relational database structure 52 stores an initial route from a source device to a destination device, as required by claim 1.

Even assuming, for the sake of argument, that the disclosure of McAllister et al. can reasonably be construed to disclose a non-alternative-enabled node that includes a storage space to store an initial route from a source device to a destination device, as required by claim 1 (a point that Appellants do not concede), Appellants submit that one skilled in the art would not have been motivated to incorporate this alleged teaching of McAllister et al. into the Haskin et al. system, absent impermissible hindsight.

With respect to motivation, the Examiner alleges "it would have been obvious ... to modify the system of Haskin with the teaching of McAllister to provide a network for forwarding packets from a source device to a destination device in order to reduce the probability of packet loss in a network" (final Office Action, pg. 7). Appellants submit that the Examiner's motivation is merely a conclusory statement regarding an alleged benefit of the combination. Such motivation statements have consistently been found to be insufficient for establishing a *prima facie* case of obviousness. Moreover, Haskin et al. already discloses a network for forwarding packets from a source device to a destination device and re-routing packets from a primary path to an alternate path to reduce the probability of packet loss (see, for example, col. 2, lines 9-50). The Examiner does not explain why incorporating the ability to store an initial route in a non-alternative-route-enabled node would enable Haskin et al.'s system to reduce the probability of packet loss. It is clear that the Examiner's motivation to combine these documents is based on impermissible hindsight.

For at least the foregoing reasons, Appellants submit that the rejection of claim 1 under 35 U.S.C. § 103(a) based on Haskin et al. and McAllister et al. is improper. Accordingly, Appellants request that the rejection of claim 1 be reversed. Moreover, since claims 2-4 and 6

depend from claim 1, Appellants further request that the rejection of these claims be reversed for at least the reasons given above with respect to claim 1.

2. Claims 8 and 12.

Claim 8 is directed to a method for forwarding packets from a source device to a destination device in a network of interconnected elements including nodes and links. The method includes determining an initial route, the initial route including at least one alternative-route enabled node and at least one non-alternative-route-enabled node, the at least one alternative-route-enabled node and the at least one non-alternative-route-enabled node storing an initial route from the source device to the destination device; determining an alternative route by identifying the at least one alternative-route-enabled node in the initial route, identifying downstream interconnected elements, and generating the alternative route based on the identified at least one alternative-route-enabled node and the identified downstream interconnected elements; forwarding packets on the initial route; detecting a failed element; and automatically forwarding packets on the alternative route without communicating with either the source or the destination. Haskin et al. and McAllister et al., whether taken alone or in any reasonable combination, do not disclose or suggest this combination of features.

For example, Haskin et al. and McAllister et al. do not disclose or suggest determining an initial route including at least one alternative-route enabled node and at least one non-alternative-route-enabled node, where the at least one alternative-route-enabled node and the at least one non-alternative-route-enabled node store the initial route from the source device to the destination device. The Examiner appears to admit that Haskin et al. does not disclose this feature of claim 8 (final Office Action, pp. 8-10). The Examiner relies on Fig. 5 and col. 13,

lines 35-67, of McAllister et al. for allegedly disclosing the above feature of claim 8 (final Office Action, pg. 10). Appellants respectfully disagree with the Examiner's interpretation of McAllister et al.

At the outset, Appellants submit that McAllister et al. does not disclose or suggest an initial route that includes at least one alternative-route-enabled node and at least one non-alternative-route-enabled node. The Examiner relies on element 32A in McAllister et al. as allegedly corresponding to an alternative-route-enabled node, on element 32B in McAllister et al. as allegedly corresponding to a non-alternative-route-enabled node (final Office Action, pg. 10). Appellants respectfully disagree with the Examiner's interpretation of McAllister et al.

Element 32A in McAllister et al. corresponds to an ingress node (see, for example, col. 8, lines 26-31) and element 32B corresponds to a node, which is attached to ingress node 32A (see, for example, Figs. 2A and 2B, and col. 6, lines 10-12). McAllister et al. in no way discloses or suggests that node 32B is a non-alternative-route-enabled node. In stark contrast, McAllister et al. specifically discloses that node 32B may transmit data via two separate paths (see, for example, col. 7, lines 35-49, where McAllister et al. specifically discloses that operator directed route (ODR) soft permanent virtual circuits (SPVCs) can be established through a primary route including node 32A; node 32B, line 36BC; node 32C; and 32D, and an alternate route including node 32A; node 32B, line 36BC'; node 32C; and 32D). Thus, node 32B is capable of transmitting data via two separate routes (a primary route and an alternative route). Therefore, node 32B would be construed as an alternative-route-enabled node and not, as alleged by the Examiner, a non-alternative-route-enabled node.

Even assuming, for the sake of argument, that one skilled in the art at the time of

Appellants' invention could reasonably construe McAllister et al.'s node 32B as a non-alternative-route-enabled node (a point that Appellants do not concede), Appellants submit that McAllister et al. does not disclose or suggest that node 32B stores an initial route from a source device to a destination device, as required by claim 8. As set forth above, the Examiner relies on Fig. 5 and col. 13, lines 35-67, of McAllister et al. for allegedly disclosing the above feature of claim 8 (final Office Action, pg. 10).

Fig. 5 of McAllister et al. depicts a relational database structure 52 that is used by a network node 32 to keep track of ODR SPVCs managed by that network node 32 (col. 13, lines 35-37). McAllister et al. appears to disclose that ingress node 32A manages the ODR SPVCs of network 30 (see, for example, col. 9, lines 6-21, where McAllister et al. specifically discloses that ingress node 32A stores the configuration data for and initiates the ODR SPVC and that the nature of the ODR SPVC is transparent to the remainder of network 30). The Examiner has not pointed to any section of McAllister et al. that discloses or suggests that node 32B stores an initial route from a source device to a destination device, as required by claim 8.

At col. 13, lines 35-67, McAllister et al. discloses:

> FIG. 5 illustrates a preferred relational database structure 52 used by a given network node 32 to keep track of ODR SPVCs managed thereby. The node database structure includes the following tables: (a) an indexed table 54 of P-NNI network node identifiers; (b) an indexed table 56 of "compressed" ODR SPVC DTLs, as explained in greater detail below, and an ODR SPVC list 58, part 59 of which is stored in random access memory. The ODR SPVC database 52 includes one record for each ODR SPVC which has originated from or is managed by the node. Each ODR SPVC record includes:
>
> (a) an "operatorDrtRtng" field 60 which specifies whether the corresponding ODR SPVC is a conventional SPVC or an ODR SPVC;
>
> (b) a "prmNtwrkPathIndex" field 62 which points to an entry in the compressed DTL table that represents the primary path of an ODR SPVC;

> (c) a "altNtwrkPathIndex" field 64 which points to the alternate path entry in the compressed DTL table for the ODR SPVC; and
>
> (d) a "re-route scheme" field 66 which stores the re-route scheme for the ODR SPVC.
>
> In this manner, the source node of an SPVC can determine whether the SPVC is an ODR SPVC, and, if so, determine the attributes associated with the ODR SPVC in order to take appropriate action in the event of link failure.
>
> The compressed DTL table 56 has a "compressed DTL" field 68 for storing network paths in a compressed format. A compressed network path is illustrated in greater detail at reference no. 68' and comprises a sequence of link identifiers (i.e. P-NNIPortIdfield 70), and pointers 72 to the node table 54. In the preferred embodiment, P-NNIPortId=0 signifies that no link has been specified by the operator, whereby.

This section of McAllister et al. discloses, as set forth above, a relational database structure 52 that is used by a network node 32 to keep track of ODR SPVCs managed by that network node 32. This section of McAllister et al. in no way discloses or suggests that the node that keeps track of an ODR SPVC is a non-alternative-route-enabled node. In fact, McAllister et al. appears to disclose that ingress node 32A (an alternative-route-enabled node) manages the ODR SPVCs of network 30 (see, for example, col. 9, lines 6-21, where McAllister et al. specifically discloses that ingress node 32A stores the configuration data for and initiates the ODR SPVC and that the nature of the ODR SPVC is transparent to the remainder of network 30). Therefore, this section of McAllister et al. cannot disclose or suggest a non-alternative-route-enabled node that stores an initial route from a source device to a destination device, as required by claim 8.

Moreover, McAllister et al. does not disclose or suggest that relational database structure 52 stores an initial route from a source device to a destination device, as required by claim 8. Instead, McAllister et al. specifically discloses that relational database structure 52 stores an

indexed table 54 of P-NNI network node identifiers, and an indexed table 56 of compressed ODR SPVC designated transit lists (DTLs) and an ODR SPVC list 58. McAllister et al. does not disclose or suggest that relational database structure 52 stores an initial route from a source device to a destination device, as required by claim 8.

Even assuming, for the sake of argument, that the disclosure of McAllister et al. can reasonably be construed to disclose a non-alternative-enabled node that stores an initial route from a source device to a destination device, as required by claim 8 (a point that Appellants do not concede), Appellants submit that one skilled in the art would not have been motivated to incorporate this alleged teaching of McAllister et al. into the Haskin et al. system, absent impermissible hindsight.

With respect to motivation, the Examiner alleges "it would have been obvious ... to modify the system of Haskin with the teaching of McAllister to provide a network for forwarding packets from a source device to a destination device in order to reduce the probability of packet loss in a network" (final Office Action, pg. 7). Appellants submit that the Examiner's motivation is merely a conclusory statement regarding an alleged benefit of the combination. Such motivation statements have consistently been found to be insufficient for establishing a *prima facie* case of obviousness. Moreover, Haskin et al. already discloses a network for forwarding packets from a source device to a destination device and re-routing packets from a primary path to an alternate path to reduce the probability of packet loss (see, for example, col. 2, lines 9-50). The Examiner does not explain why incorporating the ability to store an initial route in a non-alternative-route-enabled node would enable Haskin et al.'s system to reduce the probability of packet loss. It is clear that the Examiner's motivation to combine these documents is based on

impermissible hindsight.

For at least the foregoing reasons, Appellants submit that the rejection of claim 8 under 35 U.S.C. § 103(a) based on Haskin et al. and McAllister et al. is improper. Accordingly, Appellants request that the rejection be reversed. Moreover, since claim 12 depends from claim 8, Appellants further request that the rejection of this claim be reversed for at least the reasons given above with respect to claim 8.

3. Claim 9.

Claim 9 depends from claim 8. Therefore, claim 9 is patentable over Haskin et al. and McAllister et al., whether taken alone or in any reasonable combination, for at least the reasons given above with respect to claim 8. Moreover, claim 9 recites additional features not disclosed or suggested by Haskin et al. and McAllister et al.

Claim 9 recites that the determining an initial route includes determining a short path from the destination device to the source device within the network, refining the path according to administrative constraints, and establishing the path as the initial route. Haskin et al. and McAllister et al., whether taken alone or in any reasonable combination, do not disclose or suggest this combination of features.

For example, Haskin et al. and McAllister et al. do not disclose or suggest refining a determined short path according to administrative constraints. The Examiner appears to admit that Haskin et al. does not disclose this feature and relies on col. 8, lines 10-18, of McAllister et al. for allegedly disclosing the above feature of claim 9 (final Office Action, pg. 20). Appellants respectfully disagree with the Examiner's interpretation of McAllister et al.

At col. 8, lines 10-18, McAllister et al. discloses:

> or scheme through the user interface means, thereby indicating how strictly an ODR SPVC is restricted to routes along the nodes and links provisioned by the operator when circumstances dictate the ODR SPVC must be re-routed. The preferred re-routing schemes, which are discussed in greater detail below, include:
>
> (a) primary path;
>
> (b) primary path-alternate path;
>
> (c) primary path-any path, and

This section of McAllister et al. discloses three different re-routing schemes: primary path, primary path-alternate path, and primary path-any path. Neither this section of McAllister et al. nor any other section of McAllister et al. discloses or suggests refining a determined short path according to administrative constraints, as required by claim 9.

For at least these additional reasons, Appellants submit that the rejection of claim 9 under 35 U.S.C. § 103(a) based on Haskin et al. and McAllister et al. is improper. Accordingly, Appellants request that the rejection be reversed.

4. Claim 11.

Claim 11 depends from claim 8. Therefore, claim 11 is patentable over Haskin et al. and McAllister et al., whether taken alone or in any reasonable combination, for at least the reasons given above with respect to claim 8. Moreover, claim 11 recites additional features not disclosed or suggested by Haskin et al. and McAllister et al.

Claim 11 recites that the determining an alternative route includes determining a shortest route from a node preceding the failed element to the destination device within the network, refining the route to exclude the failed element on the initial route, and establishing the alternative route for forwarding packets. Haskin et al. and McAllister et al., whether taken alone or in any reasonable combination, do not disclose or suggest this combination of features.

For example, Haskin et al. and McAllister et al. do not disclose or suggest determining a shortest route from a node preceding the failed element to the destination device in the network. The Examiner appears to admit that Haskin et al. does not disclose this feature and relies on col. 8, lines 5-8, of McAllister et al. for allegedly disclosing the above feature of claim 11 (final Office Action, pg. 20). Appellants respectfully disagree with the Examiner's interpretation of McAllister et al.

At col. 8, lines 5-8, McAllister et al. discloses:

> the shortest path or least cost. This path is displayed, either textually or graphically, to the operator, who may then confirm or edit the path chosen by the software running on the NMS 46.

This section of McAllister et al. discloses that a shortest path or least cost path is chosen by software running on NMS 46 and can be confirmed or edited by an operator. This section of McAllister et al. in no way discloses or suggests determining a shortest route from a node preceding a failed element to the destination device in the network, as required by claim 11.

For at least these additional reasons, Appellants submit that the rejection of claim 11 under 35 U.S.C. § 103(a) based on Haskin et al. and McAllister et al. is improper. Accordingly, Appellants request that the rejection be reversed.

5. Claim 14.

Independent claim 14 is directed to a method for forwarding packets from a source device to a destination device in a network of interconnected elements including nodes and links. The method includes determining an initial route by determining a short path from the destination device to the source device within the network, refining the path according to administrative constraints, and establishing the path as the initial route, the initial route being prioritized to

establish a hierarchy for preemption in routing network traffic; determining an alternative route; forwarding packets on the initial route; detecting a failed element; and automatically forwarding packets on the alternative route without communicating with either the source or the destination. Haskin et al. and McAllister et al., whether taken alone or in any reasonable combination, do not disclose or suggest this combination of features.

For example, Haskin et al. and McAllister et al. do not disclose or suggest determining an initial route by determining a short path from the destination device to the source device within the network, refining the path according to administrative constraints, and establishing the path as the initial route, the initial route being prioritized to establish a hierarchy for preemption in routing network traffic. The Examiner does not address these features in the final Office Action. Instead, the Examiner appears to have simply copied the rejection of claim 1 (see final Office Action, pp. 11-13). Claim 1, however, recites different features than are recited in claim 14. For example, claim 1 does not recite that an initial route is prioritized to establish a hierarchy for preemption in routing network traffic, as required by claim 14. The Examiner has not established a *prima facie* case of obviousness with respect to claim 14.

With respect to claim 9, the Examiner relies on col. 8, lines 10-18, of McAllister et al. for allegedly disclosing refining a determined short path according to administrative constraints (final Office Action, pg. 20). Appellants respectfully disagree with the Examiner's interpretation of McAllister et al.

At col. 8, lines 10-18, McAllister et al. discloses:

> or scheme through the user interface means, thereby indicating how strictly an ODR SPVC is restricted to routes along the nodes and links provisioned by the operator when circumstances dictate the ODR SPVC must be re-routed. The preferred re-routing schemes, which are discussed in greater detail below, include:

(a) primary path;

(b) primary path-alternate path;

(c) primary path-any path, and

This section of McAllister et al. discloses three different re-routing schemes: primary path, primary path-alternate path, and primary path-any path. Neither this section of McAllister et al. nor any other section of McAllister et al. discloses or suggests refining a determined short path according to administrative constraints, as required by claim 14.

For at least the foregoing reasons, Appellants submit that the rejection of claim 14 under 35 U.S.C. § 103(a) based on Haskin et al. and McAllister et al. is improper. Accordingly, Appellants request that the rejection be reversed.

6. Claims 18 and 21.

Independent claim 18 is directed to a method for locally rerouting packets traveling on an established route when a node in a network of interconnected nodes fails. The method includes computing, at select intermediary nodes along the established route, an alternative route leading from the select intermediary node to the destination device of the established route; storing, at each of the select intermediary nodes, the alternative route; determining locally that the established route has failed; and automatically forwarding packets on the alternative route. Haskin et al. and McAllister et al., whether taken alone or in any reasonable combination, do not disclose or suggest this combination of features.

For example, Haskin et al. and McAllister et al. do not disclose or suggest computing, at select intermediary nodes along the established route, an alternative route leading from the select intermediary node to the destination device of the established route and storing, at each of the

select intermediary nodes, the alternative route. The Examiner does not address these features in the final Office Action. Instead, the Examiner appears to have simply copied the rejection of claim 1 (see final Office Action, pp. 14-16). Claim 1, however, recites different features than are recited in claim 18. For example, claim 1 does not recite computing, at select intermediary nodes along the established route, an alternative route leading from the select intermediary node to the destination device of the established route and storing, at each of the select intermediary nodes, the alternative route, as required by claim 18. The Examiner has not established a *prima facie* case of obviousness with respect to claim 18.

Nonetheless, Haskin et al. and McAllister et al. do not disclose or suggest computing, at select intermediary nodes along the established route, an alternative route leading from the select intermediary node to the destination device of the established route and storing, at each of the select intermediary nodes, the alternative route. Haskin et al. discloses establishing an alternative opposite direction unidirectional label switched path between a last hop switch 5 and a destination switch 7 (col. 3, line 61, to col. 4, line 45). Haskin et al. in no way discloses or suggests computing, at select intermediary nodes along the established route, an alternative route leading from the select intermediary node to the destination device of the established route and storing, at each of the select intermediary nodes, the alternative route, as required by claim 18.

McAllister et al. discloses that alternate routes are provided by an operator of network management system 46 (col. 7, line 50, to col. 8, line 8). McAllister et al. does not disclose or suggest computing, at select intermediary nodes along the established route, an alternative route leading from the select intermediary node to the destination device of the established route and storing, at each of the select intermediary nodes, the alternative route, as required by claim 18.

For at least the foregoing reasons, Appellants submit that the rejection of claim 18 under 35 U.S.C. § 103(a) based on Haskin et al. and McAllister et al. is improper. Accordingly, Appellants request that the rejection be reversed. Moreover, since claim 21 depends from claim 18, Appellants further request that the rejection of this claim be reversed for at least the reasons given above with respect to claim 18.

7. Claim 24.

Independent claim 24 is directed to a network for forwarding packets from a source device to a destination device and including a plurality of intermediate network nodes. The plurality of intermediate network nodes includes at least one first node configured to store an initial route from the source device to the destination device and at least one alternative route from the source device to the destination device, detect a failure in a downstream network node in the initial route, and automatically forward a packet to a node on one of the at least one alternative route in response to detecting the failure; and at least one second node configured to store the initial route, detect a failure in a downstream network node in the initial route, and forward a failure message to an upstream first node in response to detecting the failure, the failure message causing the upstream first node to automatically forward a packet to a node on one of the at least one alternative route. Haskin et al. and McAllister et al., whether taken alone or in any reasonable combination, do not disclose or suggest this combination of features.

For example, Haskin et al. and McAllister et al. do not disclose or suggest at least one first node and at least one second node that store an initial route from a source device to a destination device. The Examiner appears to admit that Haskin et al. does not disclose this feature and relies on Fig. 5 and col. 13, lines 35-67, of McAllister et al. for allegedly disclosing

the above feature of claim 24 (final Office Action, pg. 19). Moreover, the Examiner appears to allege that McAllister et al.'s node 32A corresponds the first node and McAllister et al.'s node 32B corresponds to the second node recited in claim 24 (final Office Action, pg. 18). Appellants respectfully disagree with the Examiner's interpretation of McAllister et al.

Fig. 5 of McAllister et al. depicts a relational database structure 52 that is used by a network node 32 to keep track of ODR SPVCs managed by that network node 32 (col. 13, lines 35-37). McAllister et al. appears to disclose that ingress node 32A manages the ODR SPVCs of network 30 (see, for example, col. 9, lines 6-21, where McAllister et al. specifically discloses that ingress node 32A stores the configuration data for and initiates the ODR SPVC and that the nature of the ODR SPVC is transparent to the remainder of network 30). The Examiner has not pointed to any section of McAllister et al. that discloses or suggests that node 32B stores an initial route from a source device to a destination device, as required by claim 24.

At col. 13, lines 35-67, McAllister et al. discloses:

> FIG. 5 illustrates a preferred relational database structure 52 used by a given network node 32 to keep track of ODR SPVCs managed thereby. The node database structure includes the following tables: (a) an indexed table 54 of P-NNI network node identifiers; (b) an indexed table 56 of "compressed" ODR SPVC DTLs, as explained in greater detail below, and an ODR SPVC list 58, part 59 of which is stored in random access memory. The ODR SPVC database 52 includes one record for each ODR SPVC which has originated from or is managed by the node. Each ODR SPVC record includes:
>
> (a) an "operatorDrtRtng" field 60 which specifies whether the corresponding ODR SPVC is a conventional SPVC or an ODR SPVC;
>
> (b) a "prmNtwrkPathIndex" field 62 which points to an entry in the compressed DTL table that represents the primary path of an ODR SPVC;
>
> (c) a "altNtwrkPathIndex" field 64 which points to the alternate path entry in the compressed DTL table for the ODR SPVC; and

> (d) a "re-route scheme" field 66 which stores the re-route scheme for the ODR SPVC.
>
> In this manner, the source node of an SPVC can determine whether the SPVC is an ODR SPVC, and, if so, determine the attributes associated with the ODR SPVC in order to take appropriate action in the event of link failure.
>
> The compressed DTL table 56 has a "compressed DTL" field 68 for storing network paths in a compressed format. A compressed network path is illustrated in greater detail at reference no. 68' and comprises a sequence of link identifiers (i.e. P-NNIPortIdfield 70), and pointers 72 to the node table 54. In the preferred embodiment, P-NNIPortId=0 signifies that no link has been specified by the operator, whereby.

This section of McAllister et al. discloses, as set forth above, a relational database structure 52 that is used by a network node 32 to keep track of ODR SPVCs managed by that network node 32. This section of McAllister et al. in no way discloses or suggests that the node that keeps track of an ODR SPVC is a second node, as required by claim 24. In fact, McAllister et al. appears to disclose that ingress node 32A (a first node) manages the ODR SPVCs of network 30 (see, for example, col. 9, lines 6-21, where McAllister et al. specifically discloses that ingress node 32A stores the configuration data for and initiates the ODR SPVC and that the nature of the ODR SPVC is transparent to the remainder of network 30). Therefore, this section of McAllister et al. cannot disclose or suggest a second node that stores an initial route from a source device to a destination device, as required by claim 24.

Moreover, McAllister et al. does not disclose or suggest that relational database structure 52 stores an initial route from a source device to a destination device, as required by claim 24. Instead, McAllister et al. specifically discloses that relational database structure 52 stores an indexed table 54 of P-NNI network node identifiers, and an indexed table 56 of compressed ODR SPVC designated transit lists (DTLs) and an ODR SPVC list 58. McAllister et al. does not

disclose or suggest that relational database structure 52 stores an initial route from a source device to a destination device, as required by claim 24.

Even assuming, for the sake of argument, that the disclosure of McAllister et al. can reasonably be construed to disclose a first node and a second node that store an initial route from a source device to a destination device, as required by claim 24 (a point that Appellants do not concede), Appellants submit that one skilled in the art would not have been motivated to incorporate this alleged teaching of McAllister et al. into the Haskin et al. system, absent impermissible hindsight.

With respect to motivation, the Examiner alleges "it would have been obvious ... to modify the system of Haskin with the teaching of McAllister to provide a network for forwarding packets from a source device to a destination device in order to reduce the probability of packet loss in a network" (final Office Action, pg. 19). Appellants submit that the Examiner's motivation is merely a conclusory statement regarding an alleged benefit of the combination. Such motivation statements have consistently been found to be insufficient for establishing a *prima facie* case of obviousness. Moreover, Haskin et al. already discloses a network for forwarding packets from a source device to a destination device and re-routing packets from a primary path to an alternate path to reduce the probability of packet loss (see, for example, col. 2, lines 9-50). The Examiner does not explain why incorporating the ability to store an initial route in a first node and a second node would enable Haskin et al.'s system to reduce the probability of packet loss. It is clear that the Examiner's motivation to combine these documents is based on impermissible hindsight.

For at least these additional reasons, Appellants submit that the rejection of claim 24

under 35 U.S.C. § 103(a) based on Haskin et al. and McAllister et al. is improper. Accordingly, Appellants request that the rejection be reversed.

B. The rejection under 35 U.S.C. § 103(a) based on Haskin et al. (U.S. Patent No. 6,813,242) in view of McAllister et al. (U.S. Patent No. 6,697,329), and further in view of Rexford (U.S. Patent No. 6,633,544) should be reversed.

1. Claim 10.

Claim 10 depends from claim 9. The disclosure of Rexford does not remedy the deficiencies in the disclosures of Haskin et al. and McAllister et al. set forth above with respect to claim 9. Therefore, claim 10 is patentable over Haskin et al., McAllister et al., and Rexford, whether taken alone or in any reasonable combination, for at least the reasons given above with respect to claim 9. Moreover, claim 10 recites another feature not disclosed or suggested by Haskin et al., McAllister et al., and Rexford.

Claim 10 recites that the refining a path comprises rejecting the path exceeding bandwidth allocation and hop limit. Haskin et al., McAllister et al., and Rexford, whether taken alone or in any reasonable combination, do not disclose or suggest this feature. The Examiner does not address this feature in the final Office Action. Accordingly, a *prima facie* case of obviousness has not been established with respect to claim 10.

For at least these additional reasons, Appellants submit that the rejection of claim 10 under 35 U.S.C. § 103(a) based on Haskin et al., McAllister et al., and Rexford is improper. Accordingly, Appellants request that the rejection be reversed.

2. Claim 13.

Claim 13 depends from claim 8. The disclosure of Rexford does not remedy the

deficiencies in the disclosures of Haskin et al. and McAllister et al. set forth above with respect to claim 8. Therefore, claim 13 is patentable over Haskin et al., McAllister et al., and Rexford, whether taken alone or in any reasonable combination, for at least the reasons given above with respect to claim 8. Moreover, claim 13 recites other features not disclosed or suggested by Haskin et al., McAllister et al., and Rexford.

Claim 13 recites that the determining the alternative route comprises reserving bandwidth available on the initial route; generating the alternative route by invoking a routing protocol; refining the alternative route by excluding the failed element; and establishing the alternative route. Haskin et al., McAllister et al., and Rexford, whether taken alone or in any reasonable combination, do not disclose or suggest this combination of features.

For example, Haskin et al., McAllister et al., and Rexford do not disclose or suggest determining an alternate route that includes reserving bandwidth available on an initial route. The Examiner admits that Haskin et al. and McAllister et al. do not disclose this feature (final Office Action, pg. 21). The Examiner relies on col. 16, lines 15-22, of Rexford for allegedly disclosing the above feature of claim 13 (final Office Action, pg. 21). Appellants respectfully disagree with the Examiner's interpretation of Rexford.

At col. 16, lines 12-22, Rexford discloses:

> To evaluate the cost-performance trade-offs of precomputed routes, an event-driven simulator that modeled link-state routing at the connection level was developed. The simulator operated by choosing a route for each incoming connection based on a throughput requirement (bandwidth b) and the available bandwidth in the network, based on the source's view of link-state information. Then, hop-by-hop signaling was employed to reserve the requested bandwidth at each link in the route. That is, if a link had a reserved bandwidth with utilization u, admitting the new connection increased the reservation to u=u+b.

This section of Rexford is directed to determining the cost-performance tradeoffs of pre-

computed routes. This section of Rexford discloses the use of hop-by-hop signaling to reserve bandwidth at each link in a route. This section of Rexford in no way relates to determining an alternate route that includes reserving bandwidth available on an initial route, as required by claim 13.

Even assuming, for the sake of argument, that that the disclosure of Rexford can reasonably be construed to disclose reserving bandwidth available on an initial route when determining an alternate route, as required by claim 13 (a point that Appellants do not concede), Appellants submit that one skilled in the art would not have been motivated to incorporate this alleged teaching of Rexford into the Haskin et al. system, absent impermissible hindsight.

With respect to motivation, the Examiner alleges "it would have been obvious … to modify the combined system (Haskin – McAllister) to check bandwidth allocation on the alternative route in order to guarantee bandwidth" (final Office Action, pg. 21). Appellants submit that the Examiner's motivation is merely a conclusory statement regarding an alleged benefit of the combination. Such motivation statements have consistently been found to be insufficient for establishing a *prima facie* case of obviousness. The Examiner does not explain why incorporating the ability to reserve bandwidth on an initial route when determining an alternate route would enable Haskin et al.'s system to guarantee bandwidth. It is clear that the Examiner's motivation to combine these documents is based on impermissible hindsight.

For at least these additional reasons, Appellants submit that the rejection of claim 13 under 35 U.S.C. § 103(a) based on Haskin et al., McAllister et al., and Rexford is improper. Accordingly, Appellants request that the rejection be reversed.

3. Claim 15.

Claim 15 depends from claim 8. The disclosure of Rexford does not remedy the deficiencies in the disclosures of Haskin et al. and McAllister et al. set forth above with respect to claim 8. Therefore, claim 15 is patentable over Haskin et al., McAllister et al., and Rexford, whether taken alone or in any reasonable combination, for at least the reasons given above with respect to claim 8. Moreover, claim 15 recites other features not disclosed or suggested by Haskin et al., McAllister et al., and Rexford.

Claim 15 recites that the determining the alternative route comprises checking bandwidth allocation. The Examiner appears to admit that Haskin et al. and McAllister et al. do not disclose this feature (final Office Action, pg. 21). The Examiner relies on col. 16, lines 15-22, of Rexford for allegedly disclosing the above feature of claim 15 (final Office Action, pg. 21). Appellants respectfully disagree with the Examiner's interpretation of Rexford.

At col. 16, lines 12-22, Rexford discloses:

> To evaluate the cost-performance trade-offs of precomputed routes, an event-driven simulator that modeled link-state routing at the connection level was developed. The simulator operated by choosing a route for each incoming connection based on a throughput requirement (bandwidth b) and the available bandwidth in the network, based on the source's view of link-state information. Then, hop-by-hop signaling was employed to reserve the requested bandwidth at each link in the route. That is, if a link had a reserved bandwidth with utilization u, admitting the new connection increased the reservation to u=u+b.

This section of Rexford is directed to determining the cost-performance tradeoffs of pre-computed routes. This section of Rexford discloses the use of hop-by-hop signaling to reserve bandwidth at each link in a route. This section of Rexford in no way relates to determining an alternate route that includes checking bandwidth allocation, as required by claim 15.

Even assuming, for the sake of argument, that that the disclosure of Rexford can reasonably be construed to disclose checking bandwidth available when determining an alternate

route, as required by claim 15 (a point that Appellants do not concede), Appellants submit that one skilled in the art would not have been motivated to incorporate this alleged teaching of Rexford into the Haskin et al. system, absent impermissible hindsight.

With respect to motivation, the Examiner alleges "it would have been obvious ... to modify the combined system (Haskin – McAllister) to check bandwidth allocation on the alternative route in order to guarantee bandwidth" (final Office Action, pg. 21). Appellants submit that the Examiner's motivation is merely a conclusory statement regarding an alleged benefit of the combination. Such motivation statements have consistently been found to be insufficient for establishing a *prima facie* case of obviousness. The Examiner does not explain why incorporating the ability to check bandwidth allocation when determining an alternate route would enable Haskin et al.'s system to guarantee bandwidth. It is clear that the Examiner's motivation to combine these documents is based on impermissible hindsight.

For at least these additional reasons, Appellants submit that the rejection of claim 15 under 35 U.S.C. § 103(a) based on Haskin et al., McAllister et al., and Rexford is improper. Accordingly, Appellants request that the rejection be reversed.

4. Claim 16.

Claim 16 depends from claim 15. Therefore, claim 16 is patentable over Haskin et al., McAllister et al., and Rexford, whether taken alone or in any reasonable combination, for at least the reasons given above with respect to claim 15. Moreover, claim 16 recites other features not disclosed or suggested by Haskin et al., McAllister et al., and Rexford.

Claim 16 recites that the checking bandwidth allocation comprises dynamically balancing capacity of nodes and links. The Examiner appears to admit that Haskin et al. and McAllister et

al. do not disclose this feature (final Office Action, pg. 21). The Examiner relies on col. 5, lines 34-35, of Rexford for allegedly disclosing the above feature of claim 16 (final Office Action, pg. 21). Appellants respectfully disagree with the Examiner's interpretation of Rexford.

At col. 5, lines 34-37, Rexford discloses:

> Path precomputation schemes benefit from having multiple candidate routes to each destination to balance network load and have additional routing choices in the case of signaling failure.

This section of Rexford discloses that by having multiple candidate routes to each destination, network load can be balanced. This section of Rexford in no way discloses or suggests that determining an alternate route includes checking bandwidth allocation, which includes dynamically balancing capacity of nodes and links, as required by claim 16.

Even assuming, for the sake of argument, that that the disclosure of Rexford can reasonably be construed to disclose checking bandwidth available when determining an alternate route, as required by claim 16 (a point that Appellants do not concede), Appellants submit that one skilled in the art would not have been motivated to incorporate this alleged teaching of Rexford into the Haskin et al. system, absent impermissible hindsight.

With respect to motivation, the Examiner alleges "it would have been obvious … to modify the combined system (Haskin – McAllister) to check bandwidth allocation on the alternative route in order to guarantee bandwidth" (final Office Action, pg. 21). Appellants submit that the Examiner's motivation is merely a conclusory statement regarding an alleged benefit of the combination. Such motivation statements have consistently been found to be insufficient for establishing a *prima facie* case of obviousness. The Examiner does not explain why incorporating the ability to dynamically balance capacity of nodes and links when

determining an alternate route would enable Haskin et al.'s system to guarantee bandwidth. It is clear that the Examiner's motivation to combine these documents is based on impermissible hindsight.

For at least these additional reasons, Appellants submit that the rejection of claim 16 under 35 U.S.C. § 103(a) based on Haskin et al., McAllister et al., and Rexford is improper. Accordingly, Appellants request that the rejection be reversed.

5. Claim 17.

Claim 17 depends from claim 8. The disclosure of Rexford does not remedy the deficiencies in the disclosures of Haskin et al. and McAllister et al. set forth above with respect to claim 8. Therefore, claim 17 is patentable over Haskin et al., McAllister et al., and Rexford, whether taken alone or in any reasonable combination, for at least the reasons given above with respect to claim 8. Moreover, claim 17 recites other features not disclosed or suggested by Haskin et al., McAllister et al., and Rexford.

Claim 17 recites that the determining the alternative route comprises reserving bandwidth available on the initial route; identifying a plurality of nodes associated with the failed element according to network configuration information; generating the alternative route excluding the failed element and the plurality of nodes; and establishing the alternative route. The Examiner appears to admit that Haskin et al. and McAllister et al. do not disclose this feature (final Office Action, pg. 21). The Examiner relies on col. 16, lines 15-22, of Rexford for allegedly disclosing the above features of claim 17 (final Office Action, pg. 21). Appellants respectfully disagree with the Examiner's interpretation of Rexford.

At col. 16, lines 12-22, Rexford discloses:

> To evaluate the cost-performance trade-offs of precomputed routes, an event-driven simulator that modeled link-state routing at the connection level was developed. The simulator operated by choosing a route for each incoming connection based on a throughput requirement (bandwidth b) and the available bandwidth in the network, based on the source's view of link-state information. Then, hop-by-hop signaling was employed to reserve the requested bandwidth at each link in the route. That is, if a link had a reserved bandwidth with utilization u, admitting the new connection increased the reservation to u=u+b.

This section of Rexford is directed to determining the cost-performance tradeoffs of pre-computed routes. This section of Rexford discloses the use of hop-by-hop signaling to reserve bandwidth at each link in a route. This section of Rexford in no way relates to determining an alternate route that includes reserving bandwidth available on the initial route; identifying a plurality of nodes associated with the failed element according to network configuration information; generating the alternative route excluding the failed element and the plurality of nodes; and establishing the alternative route, as required by claim 17.

Even assuming, for the sake of argument, that that the disclosure of Rexford can reasonably be construed to disclose checking bandwidth available when determining an alternate route, as required by claim 17 (a point that Appellants do not concede), Appellants submit that one skilled in the art would not have been motivated to incorporate this alleged teaching of Rexford into the Haskin et al. system, absent impermissible hindsight.

With respect to motivation, the Examiner alleges "it would have been obvious ... to modify the combined system (Haskin – McAllister) to check bandwidth allocation on the alternative route in order to guarantee bandwidth" (final Office Action, pg. 21). Appellants submit that the Examiner's motivation is merely a conclusory statement regarding an alleged benefit of the combination. Such motivation statements have consistently been found to be insufficient for establishing a *prima facie* case of obviousness. The Examiner does not explain

why incorporating the ability to reserve bandwidth available on the initial route; identify a plurality of nodes associated with the failed element according to network configuration information; generate the alternative route excluding the failed element and the plurality of nodes; and establish the alternative route when determining an alternate route would enable Haskin et al.'s system to guarantee bandwidth. It is clear that the Examiner's motivation to combine these documents is based on impermissible hindsight.

For at least these additional reasons, Appellants submit that the rejection of claim 17 under 35 U.S.C. § 103(a) based on Haskin et al., McAllister et al., and Rexford is improper. Accordingly, Appellants request that the rejection be reversed.

6. Claim 19.

Claim 19 depends from claim 18. The disclosure of Rexford does not remedy the deficiencies in the disclosures of Haskin et al. and McAllister et al. set forth above with respect to claim 18. Therefore, claim 19 is patentable over Haskin et al., McAllister et al., and Rexford, whether taken alone or in any reasonable combination, for at least the reasons given above with respect to claim 18. Moreover, claim 19 recites other features not disclosed or suggested by Haskin et al., McAllister et al., and Rexford.

Claim 19 recites that computing the alternate route includes reserving bandwidth available on the established route. The Examiner admits that Haskin et al. and McAllister et al. do not disclose this feature (final Office Action, pg. 21). The Examiner relies on col. 16, lines 15-22, of Rexford for allegedly disclosing the above feature of claim 19 (final Office Action, pg. 21). Appellants respectfully disagree with the Examiner's interpretation of Rexford.

At col. 16, lines 12-22, Rexford discloses:

> To evaluate the cost-performance trade-offs of precomputed routes, an event-driven simulator that modeled link-state routing at the connection level was developed. The simulator operated by choosing a route for each incoming connection based on a throughput requirement (bandwidth b) and the available bandwidth in the network, based on the source's view of link-state information. Then, hop-by-hop signaling was employed to reserve the requested bandwidth at each link in the route. That is, if a link had a reserved bandwidth with utilization u, admitting the new connection increased the reservation to u=u+b.

This section of Rexford is directed to determining the cost-performance tradeoffs of pre-computed routes. This section of Rexford discloses the use of hop-by-hop signaling to reserve bandwidth at each link in a route. This section of Rexford in no way relates to computing an alternate route that includes reserving bandwidth available on an established route, as required by claim 19.

Even assuming, for the sake of argument, that that the disclosure of Rexford can reasonably be construed to disclose reserving bandwidth available on an established route when computing an alternate route, as required by claim 19 (a point that Appellants do not concede), Appellants submit that one skilled in the art would not have been motivated to incorporate this alleged teaching of Rexford into the Haskin et al. system, absent impermissible hindsight.

With respect to motivation, the Examiner alleges "it would have been obvious ... to modify the combined system (Haskin – McAllister) to check bandwidth allocation on the alternative route in order to guarantee bandwidth" (final Office Action, pg. 21). Appellants submit that the Examiner's motivation is merely a conclusory statement regarding an alleged benefit of the combination. Such motivation statements have consistently been found to be insufficient for establishing a *prima facie* case of obviousness. The Examiner does not explain why incorporating the ability to reserve bandwidth on an established route when computing an alternate route would enable Haskin et al.'s system to guarantee bandwidth. It is clear that the

Examiner's motivation to combine these documents is based on impermissible hindsight.

For at least these additional reasons, Appellants submit that the rejection of claim 19 under 35 U.S.C. § 103(a) based on Haskin et al., McAllister et al., and Rexford is improper. Accordingly, Appellants request that the rejection be reversed.

7. Claim 20.

Claim 20 depends from claim 19. The disclosure of Rexford does not remedy the deficiencies in the disclosures of Haskin et al. and McAllister et al. set forth above with respect to claim 19. Therefore, claim 20 is patentable over Haskin et al., McAllister et al., and Rexford, whether taken alone or in any reasonable combination, for at least the reasons given above with respect to claim 19. Moreover, claim 20 recites other features not disclosed or suggested by Haskin et al., McAllister et al., and Rexford.

Claim 20 recites that computing the alternate route includes locating a set of established routes with a same destination device and same administrative constraints as the established route; finding a common node, downstream from the failed node, after which the set of established routes and the established route utilize the same network elements; establishing a new route from the common node to the destination device; and incorporating the new route into the alternative route. Haskin et al., McAllister et al., and Rexford, whether taken alone or in any reasonable combination, do not disclose or suggest this combination of features. The Examiner does not address these features in the final Office Action. Accordingly, a *prima facie* case of obviousness has not been established with respect to claim 20.

For at least these additional reasons, Appellants submit that the rejection of claim 20 under 35 U.S.C. § 103(a) based on Haskin et al., McAllister et al., and Rexford is improper.

Accordingly, Appellants request that the rejection be reversed.

8. Claim 21.

Claim 21 depends from claim 18. The disclosure of Rexford does not remedy the deficiencies in the disclosures of Haskin et al. and McAllister et al. set forth above with respect to claim 18. Therefore, claim 21 is patentable over Haskin et al., McAllister et al., and Rexford, whether taken alone or in any reasonable combination, for at least the reasons given above with respect to claim 18.

VIII. CONCLUSION

In view of the foregoing arguments, Appellants respectfully solicit the Honorable Board to reverse the Examiner's rejections of claims 1-4, 6, 8-21, and 24 under 35 U.S.C. § 103.

To the extent necessary, a petition for an extension of time under 37 C.F.R. § 1.136 is hereby made. Please charge any shortage in fees due in connection with the filing of this paper, including extension of time fees, to Deposit Account No. 50-1070 and please credit any excess fees to such deposit account.

Respectfully submitted,

HARRITY SNYDER, L.L.P.

By: /John E. Harrity/
John E. Harrity
Registration No. 43,367

Date: May 8, 2006

11350 Random Hills Road
Suite 600
Fairfax, Virginia 22030
(571) 432-0800

IX. CLAIM APPENDIX

1. A network for forwarding packets from a source device to a destination device, said network including a plurality of network elements including a plurality of nodes and connecting links, the plurality of nodes including at least one alternative-route-enabled node and at least one non-alternative-route-enabled node, wherein the at least one non-alternative-route-enabled node comprises:

a storage space to store an initial route from the source device to the destination device;

a mechanism to detect failure in a downstream network element in the initial route; and

a forwarder to automatically forward a failure message upstream along the initial route to an alternative-route-enabled node, the failure message causing the alternative-route-enabled node to begin forwarding packets on an alternative route.

2. The node in claim 1, wherein the network is a connection-oriented network with a plurality of established initial routes.

3. The node in claim 2, wherein the plurality of nodes includes a label-switched router.

4. The node in claim 1, wherein the alternative route does not include the downstream network element in the initial route.

6. The node in claim 1, wherein the mechanism to detect failure sends communication packets to downstream nodes at regular intervals.

8. A method for forwarding packets from a source device to a destination device in a network of interconnected elements including nodes and links, comprising:

determining an initial route, the initial route including at least one alternative-route enabled node and at least one non-alternative-route-enabled node, the at least one alternative-route-enabled node and the at least one non-alternative-route-enabled node storing an initial route from the source device to the destination device;

determining an alternative route by identifying the at least one alternative-route-enabled node in the initial route, identifying downstream interconnected elements, and generating the alternative route based on the identified at least one alternative-route-enabled node and the identified downstream interconnected elements;

forwarding packets on the initial route;

detecting a failed element; and

automatically forwarding packets on the alternative route without communicating with either the source or the destination.

9. The method of claim 8, wherein determining the initial route further comprises:

determining a short path from the destination device to the source device within the network;

refining the path according to administrative constraints; and

establishing the path as the initial route.

10. The method of claim 9, wherein refining the path comprises rejecting the path exceeding bandwidth allocation and hop limit.

11. The method of claim 8, wherein determining the alternative route further comprises:

determining a shortest route from a node preceding the failed element to the destination device within the network;

refining the route to exclude the failed element on the initial route; and

establishing the alternative route for forwarding packets.

12. The method of claim 8, wherein detecting a failure is conducted locally by a node preceding the failed element without requiring notification of a master server or an ingress node.

13. The method of claim 8, wherein determining the alternative route comprises:

reserving bandwidth available on the initial route;

generating the alternative route by invoking a routing protocol;

refining the alternative route by excluding the failed element; and

establishing the alternative route.

14. A method for forwarding packets from a source device to a destination device in a network of interconnected elements including nodes and links, comprising:

determining an initial route by determining a short path from the destination device to the source device within the network, refining the path according to administrative constraints, and establishing the path as the initial route, the initial route being prioritized to establish a hierarchy for preemption in routing network traffic;

determining an alternative route;

forwarding packets on the initial route;

detecting a failed element; and

automatically forwarding packets on the alternative route without communicating with either the source or the destination.

15. The method of claim 8, wherein the determining the alternative route comprises checking bandwidth allocation.

16. The method of claim 15, wherein checking bandwidth allocation comprises dynamically balancing capacity of nodes and links.

17. The method of claim 8, wherein determining the alternative route comprises:

reserving bandwidth available on the initial route;

identifying a plurality of nodes associated with the failed element according to network configuration information;

generating the alternative route excluding the failed element and the plurality of nodes;

establishing the alternative route.

18. A method for locally rerouting packets traveling on an established route when a node in a network of interconnected nodes fails, the method comprising:

computing, at select intermediary nodes along the established route, an alternative route leading from the select intermediary node to the destination device of the established route;

storing, at each of the select intermediary nodes, the alternative route;

determining locally that the established route has failed; and

automatically forwarding packets on the alternative route.

19. The method of claim 18, wherein computing the alternative route comprises:

reserving bandwidth available on the established route;

identifying a plurality of nodes associated with the failed node according to network configuration information;

generating the alternative route excluding the failed node and the plurality of nodes; and

establishing the alternative route.

20. The method of claim 19, wherein computing the alternative route further comprises:

locating a set of established routes with a same destination device and same administrative constraints as the established route;

finding a common node, downstream from the failed node, after which the set of established routes and the established route utilize the same network elements;

establishing a new route from the common node to the destination device; and

incorporating the new route into the alternative route.

21. The method of claim 18, wherein determining locally that the established route has failed is conducted by a signaling protocol.

24. A network for forwarding packets from a source device to a destination device and including a plurality of intermediate network nodes, the plurality of intermediate network nodes comprising:

at least one first node configured to:

store an initial route from the source device to the destination device and at least one alternative route from the source device to the destination device,

detect a failure in a downstream network node in the initial route, and

automatically forward a packet to a node on one of the at least one alternative route in response to detecting the failure; and

at least one second node configured to:

store the initial route,

detect a failure in a downstream network node in the initial route, and

forward a failure message to an upstream first node in response to detecting the failure, the failure message causing the upstream first node to automatically forward a packet to a node on one of the at least one alternative route.

X. EVIDENCE APPENDIX

None.

XI. RELATED PROCEEDINGS APPENDIX

None.